

Stable Marriage and Roommates problems with restricted edges: Complexity and approximability[☆]

Ágnes Cseh^a, David F. Manlove^b

^a*Institute for Mathematics, Technische Universität Berlin, Sekr. MA 5-2, Straße des 17. Juni 136, 10623 Berlin, Germany*

^b*School of Computing Science, Sir Alwyn Williams Building, University of Glasgow, Glasgow G12 8QQ, UK*

Abstract

In the Stable Marriage and Roommates problems, a set of agents is given, each of them having a strictly ordered preference list over some or all of the other agents. A matching is a set of disjoint pairs of mutually acceptable agents. If any two agents mutually prefer each other to their partner, then they block the matching, otherwise, the matching is said to be stable. We investigate the complexity of finding a solution satisfying additional constraints on restricted pairs of agents. Restricted pairs can be either *forced* or *forbidden*. A stable solution must contain all of the forced pairs, while it must contain none of the forbidden pairs.

Dias et al. [7] gave a polynomial-time algorithm to decide whether such a solution exists in the presence of restricted edges. If the answer is no, one might look for a solution close to optimal. Since optimality in this context means that the matching is stable and satisfies all constraints on restricted pairs, there are two ways of relaxing the constraints by permitting a solution

[☆]A preliminary version of this paper appeared in the Proceedings of SAGT 2015: the 8th International Symposium on Algorithmic Game Theory.

Email addresses: cseh@math.tu-berlin.de (Ágnes Cseh), David.Manlove@glasgow.ac.uk (David F. Manlove)

¹Supported by COST Action IC1205 on Computational Social Choice and by the Deutsche Telekom Stiftung. Part of this work was carried out whilst visiting the University of Glasgow. Present address: School of Computer Science, Reykjavik University, Menntavegur 1, 101 Reykjavik, Iceland.

²Supported by Engineering and Physical Sciences Research Council grant EP/K010042/1.

to: (1) be blocked by as few as possible pairs, or (2) violate as few as possible constraints on restricted pairs.

Our main theorems prove that for the (bipartite) Stable Marriage problem, case (1) leads to NP-hardness and inapproximability results, whilst case (2) can be solved in polynomial time. For non-bipartite Stable Roommates instances, case (2) yields an NP-hard problem. In the case of NP-hard problems, we also discuss polynomially solvable special cases, arising from restrictions on the lengths of the preference lists, or upper bounds on the numbers of restricted pairs.

Keywords: stable matching, restricted edge, approximation algorithm
2000 MSC: 05C70, 68W40, 05C85

1. Introduction

In the classical *Stable Marriage problem* (SM) [12], a bipartite graph is given, where one colour class symbolises a set of men U and the other colour class stands for a set of women W . Man u and woman w are connected by edge uw if they find one another mutually acceptable. Each participant provides a strictly ordered preference list of the acceptable agents of the opposite gender. An edge uw *blocks* matching M if it is not in M , but each of u and w is either unmatched or prefers the other to their partner. A *stable matching* is a matching not blocked by any edge. From the seminal paper of Gale and Shapley [12], we know that the existence of such a stable solution is guaranteed and one can be found in linear time. Moreover, the solutions form a distributive lattice [24]. The two extreme points of this lattice are called the *man-* and *woman-optimal stable matchings* [12]. These assign each man/woman their best partner reachable in any stable matching. Another interesting and useful property of stable solutions is the so-called Rural Hospitals Theorem. Part of this theorem states that if an agent is unmatched in one stable matching, then all stable solutions leave him unmatched [13].

One of the most widely studied extensions of SM is the *Stable Roommates problem* (SR) [12, 18], defined on general graphs instead of bipartite graphs. The notion of a blocking edge is as defined above (except that it can now involve any two agents in general), but several results do not carry over to this setting. For instance, the existence of a stable solution is not guaranteed any more. On the other hand, there is a linear-time algorithm to find a stable matching or report that none exists [18]. Moreover, the corresponding variant

of the Rural Hospitals Theorem holds in the roommates case as well: the set of matched agents is the same for all stable solutions [16]. We summarise this observation as follows:

Theorem 1.1 (Gusfield and Irving [16]). *Given an instance of SR, the same set of agents is matched in all stable matchings.*

Both SM and SR are widely used in various applications. In markets where the goal is to maximise social welfare instead of profit, the notion of stability is especially suitable as an optimality criterion [27]. For SM, the oldest and most common area of applications is employer allocation markets [29]. On one side, job applicants are represented, while the job openings form the other side. Each application corresponds to an edge in the bipartite graph. The employers rank all applicants to a specific job offer and similarly, each applicant sets up a preference list of jobs. Given a proposed matching M of applicants to jobs, if an employer–applicant pair exists such that the position is not filled or a worse applicant is assigned to it, and the applicant received no contract or a worse contract, then this pair blocks M . In this case the employer and applicant find it mutually beneficial to enter into a contract outside of M , undermining its integrity. If no such blocking pair exists, then M is stable. Stability as an underlying concept is also used to allocate graduating medical students to hospitals in many countries [28]. SR on the other hand has applications in the area of P2P networks [11].

Forced and forbidden edges in SM and SR open the way to formulate various special requirements on the sought solution. Such edges now form part of the extended problem instance: if an edge is *forced*, it must belong to a constructed stable matching, whilst if an edge is *forbidden*, it must not. In certain market situations, a contract is for some reason particularly important, or to the contrary, not wished by the majority of the community or by the central authority in control. In such cases, forcing or forbidding the edge and then seeking a stable solution ensures that the wishes on these specific contracts are fulfilled while stability is guaranteed. Henceforth, the term *restricted edge* will be used to refer either to a forbidden edge or a forced edge. The remaining edges of the graph are referred as *unrestricted edges*.

Note that simply deleting forbidden edges or fixing forced edges and searching for a stable matching on the remaining instance does not solve the problem of finding a stable matching with restricted edges. Deleted edges (corresponding to forbidden edges, or those adjacent to forced edges)

can block that matching. Therefore, to meet both requirements on restricted edges and stability, more sophisticated methods are needed.

The attention of the community was drawn very early to the characterization of stable matchings that must contain a prescribed set of edges. In the seminal book of Knuth [24], forced edges first appeared under the term *arranged marriages*. Knuth presented an algorithm that finds a stable matching with a given set of forced edges or reports that none exists, given an instance of SM. This method runs in $O(n^2)$ time, where n denotes the number of vertices in the graph. Gusfield and Irving [16] provided an algorithm for SM based on rotations that terminates in $O(|Q|^2)$ time, following $O(n^4)$ pre-processing time, where Q is the set of forced edges. This latter method is favoured over Knuth's if multiple forced sets of small cardinality are proposed.

Forbidden edges appeared only in 2003 in the literature, and were first studied by Dias et al. [7]. In their paper, complete bipartite graphs were considered, but the methods can easily be extended to incomplete preference lists. Their main result was the following (in the following theorem, and henceforth, m is the total number of edges in the graph).

Theorem 1.2 (Dias et al. [7]). *The problem of finding a stable matching in a SM instance with forced and forbidden edges or reporting that none exists is solvable in $O(m)$ time.*

While Knuth's method relies on basic combinatorial properties of stable matchings, the other two algorithms make use of *rotations*. We refer the reader to [16] for background on these. The problem of finding a stable matching with forced and forbidden edges in an SM instance can easily be formulated as a weighted stable matching problem (that is, we seek a stable matching with minimum weight, where the weight of a matching M is the sum of the weights of the edges in M). Let us assign all forced edges weight -1 , all forbidden edges weight 1 , and all remaining edges weight 0 . A stable matching satisfying all constraints on restricted edges exists if and only if there is a stable matching of weight $-|Q|$ in the weighted instance, where Q is the set of forced edges. With the help of rotations, minimum weight stable matchings can be found in polynomial time [19, 8, 30, 9] (see the final paragraph of Section 2 for more detail on the role played by each of these references).

Since finding a weight-minimal stable matching in SR instances is an NP-hard task [8], it follows that solving the problem with forced and forbidden

edges requires different methods from the aforementioned weighted transformation. Fleiner et al. [10] showed that any SR instance with forbidden edges can be converted into another stable matching problem involving ties that can be solved in $O(m)$ time [20] and the transformation has the same time complexity as well. Forced edges can easily be eliminated by forbidding all edges adjacent to them, therefore we can state the following result.

Theorem 1.3 (Fleiner et al. [10]). *The problem of finding a stable matching in an SR instance with forced and forbidden edges or reporting that none exists is solvable in $O(m)$ time.*

As we have seen so far, answering the question as to whether a stable solution containing all forced and avoiding all forbidden edges exists can be solved efficiently in the case of both SM and SR. We thus concentrate on cases where the answer to this question is no. What kind of approximate solutions exist then and how can we find them?

Our contribution. Since optimality is defined by two criteria, it is straightforward to define approximation from those two points of view. In case BP, all constraints on restricted edges must be satisfied, and we seek a matching with the minimum number of blocking edges. In case CV, we seek a stable matching that violates the fewest constraints on restricted edges. The optimisation problems that arise from each of these cases are defined formally in Section 2.

In Section 3, we consider case BP: that is, all constraints on restricted edges must be fulfilled, while the number of blocking edges is minimised. We show that in the SM case, this problem is computationally hard and not approximable within $n^{1-\varepsilon}$ for any $\varepsilon > 0$, unless $\mathbf{P} = \mathbf{NP}$. We also discuss special cases for which this problem becomes tractable. This occurs if the maximum degree of the graph is at most 2 or if the number of blocking edges in the optimal solution is a constant. We point out a striking difference in the complexity of the two cases with only forbidden and only forced edges: the problem is polynomially solvable if the number of forbidden edges is a constant, but by contrast it is \mathbf{NP} -hard even if the instance contains a single forced edge. We also prove that when the restricted edges are either all forced or all forbidden, the optimisation problem remains \mathbf{NP} -hard even on very sparse instances, where the maximum degree of a vertex is 3.

Case CV, where the number of violated constraints on restricted edges is minimised while stability is preserved, is studied in Section 4. It is a

rather straightforward observation that in SM, the setting can be modelled and efficiently solved with the help of edge weights. Here we show that on non-bipartite graphs, the problem becomes **NP**-hard. As in case BP, we also discuss the complexity of degree-constrained restrictions and establish that the **NP**-hardness results remain intact even for graphs with degree at most 3, while the case with degree at most 2 is polynomially solvable.

A structured overview of our results for general SM and SR instances is contained in Table 1.

	Stable Marriage	Stable Roommates
case BP: min # blocking edges	NP -hard to approximate within $n^{1-\varepsilon}$	NP -hard to approximate within $n^{1-\varepsilon}$
case CV: min # violated restricted edge constraints	solvable in polynomial time	NP -hard

Table 1: Summary of results

2. Preliminaries and techniques

In this section, we introduce the notation used in the remainder of the paper and also define the key problems that we investigate later. An instance $\mathcal{I} = (G, O)$ of the *Stable Marriage problem* (SM) consists of a bipartite graph $G = (U \cup W, E)$ with n vertices and m edges, and a set O : the set of strictly ordered, but not necessarily complete preference lists. These lists are provided on the set of adjacent vertices at each vertex. The *Stable Roommates problem* (SR) differs from SM in one sense: the underlying graph G need not be bipartite. In both SM and SR, a matching M in G is sought, assigning each agent to at most one partner. If a vertex $v \in V(G)$ is matched in M , we denote by $M(v)$ the partner of v in M . An edge $uw \in E \setminus M$ *blocks* M , or forms a *blocking pair* of M if either u is unmatched or prefers w to $M(u)$, and either w is unmatched or prefers u to $M(w)$. A matching that is not blocked by any edge is called *stable*.

As already mentioned in the Introduction, an SR instance need not admit a stable solution. The number of blocking edges is a characteristic property of every matching. The set of edges blocking M is denoted by $bp(M)$. A natural goal is to find a matching minimising $|bp(M)|$; following the consensus in the literature, such a matching is called *almost stable*. This approach has a broad

literature: almost stable matchings have been investigated in SM [23, 17, 6] and SR [1, 5] instances.

All problems investigated in this paper deal with at least one set of restricted edges. The set of forbidden edges is denoted by P , while Q stands for the set of forced edges. We assume throughout the paper that $P \cap Q = \emptyset$. A matching M satisfies all constraints on restricted edges if $M \cap P = \emptyset$ and $Q \subseteq M$.

In Figure 1, a sample SM instance on four men and four women can be seen. The preference ordering is shown on the edges. For instance, vertex u_2 ranks w_1 best, then w_4 , and w_2 last. The set of forbidden edges $P = \{u_2w_2, u_3w_3\}$ is marked by dotted grey edges. The unique stable matching $M = \{u_1w_1, u_2w_2, u_3w_3, u_4w_4\}$ contains both forbidden edges. Later on, we will return to this sample instance to demonstrate approximation concepts on it.




Figure 1: A sample stable marriage instance with forbidden edges

The first approximation concept (case BP described in Section 1) is to seek a matching M that satisfies all constraints on restricted edges, but among these matchings, it admits the minimum number of blocking edges. This leads to the following problem definition.

Problem 2.1. MIN BP SR RESTRICTED

Input: $\mathcal{I} = (G, O, P, Q)$ comprising an SR instance (G, O) , a set of forbidden edges P and a set of forced edges Q .

Output: A matching M such that $M \cap P = \emptyset$, $Q \subseteq M$ and $|bp(M)| \leq |bp(M')|$ for every matching M' in G satisfying $M' \cap P = \emptyset$, $Q \subseteq M'$.

Special attention is given to two special cases of MIN BP SR RESTRICTED: in MIN BP SR FORBIDDEN, $Q = \emptyset$, while in MIN BP SR FORCED, $P = \emptyset$.

Note that an instance of MIN BP SR FORCED or MIN BP SR RESTRICTED can always be transformed into an instance of MIN BP SR FORBIDDEN by forbidding all edges that are adjacent to a forced edge. This transformation does not affect the number of blocking edges.

According to the other intuitive approximation concept (case CV described in Section 1), stability constraints need to be fulfilled, while some of the constraints on restricted edges are relaxed. The goal is to find a stable matching that violates as few constraints on restricted edges as possible.

Problem 2.2. SR MIN RESTRICTED VIOLATIONS

Input: $\mathcal{I} = (G, O, P, Q)$ comprising an SR instance (G, O) , a set of forbidden edges P and a set of forced edges Q .

Output: A stable matching M such that $|M \cap P| + |Q \setminus M| \leq |M' \cap P| + |Q \setminus M'|$ for every stable matching M' in G .

Just as in the previous approximation concept (referred to as case BP in Section 1), we separate the two subcases with only forbidden and only forced edges. If $Q = \emptyset$, SR MIN RESTRICTED VIOLATIONS is referred as SR MIN FORBIDDEN, while if $P = \emptyset$, the problem becomes SR MAX FORCED. If $P = \emptyset$ or $Q = \emptyset$ then that set is omitted from an instance of MIN BP SR RESTRICTED or SR MIN RESTRICTED VIOLATIONS as appropriate.

When considering the decision versions of the problems defined in this section, we append DEC to the problem name and add a positive integer K to the problem instance. The problem is then to decide whether a feasible solution exists with measure at most K . For example, in the case of the optimisation problem MIN BP SR FORBIDDEN, an instance of the decision problem MIN BP SR FORBIDDEN DEC comprises a tuple (G, O, P, K) , where (G, O, P) is as per the definition of MIN BP SR FORBIDDEN and K is a positive integer. The question is whether there is a matching M such that $|M \cap P| = \emptyset$ and $|bp(M)| \leq K$.

In all discussed problems, n is the number of vertices and m is the number of edges in the graph underlying the particular problem instance. When considering the restriction of any of the above problems to the case of a bipartite graph SR is replaced by SM in the problem name.

In case BP, the subcase with only forced edges can be transformed into the other subcase, simply by forbidding edges adjacent to forced edges. This straightforward transformation is not valid for case CV. Suppose a forced edge was replaced by an unrestricted edge, but all of its adjacent edges were forbidden. A solution that does not contain the original forced edge might

contain two of the forbidden edges, violating more constraints than the original solution. Yet most of our proofs are presented for the problem with only forbidden edges, and they require only slight modifications for the case with forced edges.

A powerful tool used in several proofs in our paper is to convert some of these problems into a weighted SM or SR problem, where the goal is to find a stable matching with the lowest total edge weight, taken over all stable matchings. Irving et al. [19] were the first to show that weighted SM can be solved in polynomial time, giving an $O(n^4 \log n)$ algorithm if the weight function is monotone in the preference ordering, non-negative and integral. Feder [8, 9] showed a method to drop the monotonicity requirement. He also presented the best known bound for the running time of an algorithm for finding a minimum weight stable matching in SM: $O(n^2 \cdot \log(\frac{K}{n^2} + 2) \cdot \min\{n, \sqrt{K}\})$, where K is the weight of an optimal solution. Redesigning the weight function to avoid the monotonicity requirement using Feder’s method can radically increase K . Fortunately, linear programming techniques allow the conditions to be dropped while retaining polynomial-time solvability. A simple and elegant formulation of the SM polytope is known [30] and using this, a minimum weight stable matching can be computed for all real-valued weight functions in polynomial time via linear programming. For weighted SR, finding an optimal matching is NP-hard, but 2-approximable with combinatorial methods, under the assumption of monotone, non-negative and integral weights [8]. With the help of LP methods, a 2-approximation can be found for every non-negative weight function that satisfies a special monotonicity constraint [33, 34].

3. Almost stable matchings with restricted edges

In this section, constraints on restricted edges must be fulfilled strictly, while the number of blocking edges is minimised. Our results are presented in three subsections, and most of the results are given for MIN BP SM RESTRICTED. Firstly, in Section 3.1, basic complexity results are discussed. In particular, we prove that the studied problem MIN BP SM RESTRICTED is in general NP-hard and very difficult to approximate. Thus, restricted cases are analysed in Section 3.2. First we assume that the number of forbidden, forced or blocking edges can be considered as a constant. Due to this assumption, two of the three problems that naturally follow from imposing these restrictions become tractable, but surprisingly, not all of them. Then,

degree-constrained cases are discussed. We show that the **NP**-hardness result for **MIN BP SM RESTRICTED** holds even for instances where each preference list is of length at most 3, while on graphs with maximum degree 2, the problem becomes tractable. Finally, in Section 3.3 we consider the problem **MIN BP SR RESTRICTED** and briefly elaborate on whether the results established for the bipartite case carry over to the **SR** case.

3.1. General complexity and approximability results

When minimising the number of blocking edges, one might think that removing the forbidden edges temporarily and then searching for a stable solution in the remaining instance leads to an optimal solution. Such a matching can only be blocked by forbidden edges, but as the upcoming example demonstrates, optimal solutions are sometimes blocked by unrestricted edges exclusively. In some instances, all almost stable solutions admit only non-forbidden blocking edges. Moreover, a man- or woman-optimal almost stable matching with forbidden edges does not always exist.

Let us recall the **SM** instance in Figure 1. In the graph with edge set $E(G) \setminus P$, a unique stable matching exists: $M = \{u_1w_1, u_4w_4\}$. Matching M is blocked by both forbidden edges in the original instance. On the other hand, matching $M_1 = \{u_1w_1, u_2w_4, u_4w_3\}$ is blocked by exactly one edge: $bp(M_1) = u_4w_4$. Similarly, matching $M_2 = \{u_1w_3, u_2w_1, u_4w_4\}$ is blocked only by u_1w_1 . Therefore, M_1 and M_2 are both solutions to **MIN BP SM FORBIDDEN** on this instance. One can easily check that M_1 and M_2 are the only matchings with the minimum number of blocking edges. They both are blocked only by unrestricted edges. Moreover, M_1 is better for u_1, w_1 and w_3 , whereas M_2 is preferred by u_2, u_4 and w_4 .

In Theorems 3.1 and 3.5 we present two results demonstrating the **NP**-hardness and inapproximability of special cases of **MIN BP SM RESTRICTED**.

Theorem 3.1. ***MIN BP SM FORBIDDEN DEC** and **MIN BP SM FORCED DEC** are **NP**-complete. The result holds even if all preference lists are complete.*

Proof. Clearly both problems belong to **NP**. We show the **NP**-hardness of both problems by giving a reduction from the following problem:

Problem 3.2. **MIN BP PSMI DEC**

Input: $\mathcal{I} = (G, O, K)$ comprising an **SM** instance (G, O) and a positive integer K .

Output: A perfect matching M such that $|bp(M)| \leq K$.

MIN BP PSMI denotes the minimisation version of MIN BP PSMI DEC, in which we seek a perfect matching with the minimum number of blocking pairs, taken over all perfect matchings in G . MIN BP PSMI-DEC is **NP**-complete, and unless $P = NP$, MIN BP PSMI is not approximable within a factor of $n^{1-\varepsilon}$, for any $\varepsilon > 0$ [6].

We firstly show **NP**-hardness of MIN BP SM FORBIDDEN DEC and then indicate how to adapt the proof to show a similar result for MIN BP SM FORCED DEC. We reduce from MIN BP PSMI DEC as mentioned above. Given an instance $\mathcal{I} = (G, O, K)$ of this problem we define an instance $\mathcal{I}' = (G', O', P, K)$ of MIN BP SM FORBIDDEN DEC as follows. Let $G = (V, E)$ where U and W are the two colour classes of G . Let $n = |V|$; then $|U| = |W| = n/2$. Let $U = \{u_1, u_2, \dots, u_{n/2}\}$ and let $W = \{w_1, w_2, \dots, w_{n/2}\}$. Add the vertices in V to G' . In addition, $K + 1$ new vertices representing women are added to G' . They are denoted by $Y = \{y_1, y_2, \dots, y_{K+1}\}$. Similarly, $K + 1$ new men $X = \{x_1, x_2, \dots, x_{K+1}\}$ are added to G' . Thus, each colour class of G' consists of $n/2 + K + 1$ vertices.

In O' the preference lists of vertices already in $V(G)$ are structured in three blocks. Each man u_i in the original instance \mathcal{I} keeps his preference list in O at the top of his new list in O' . After these vertices, the entire set of newly-introduced women in Y follows, in arbitrary order. Finally, the rest of the women in W , not already in u_i 's list follow, in arbitrary order. A similar ordering is used when defining the preference list of each w_j . The original list in O is followed by the vertices in X , and then the rest of the men in U follow.

The added newly-added vertices in $X \cup Y$ have different preference orderings. Man x_i 's list consists of the women in W in arbitrary order, followed by y_i , and then the women in $Y \setminus \{y_i\}$ in arbitrary order. Similarly y_j ranks all men in U first in arbitrary order, followed by x_j , and then the men in $X \setminus \{x_j\}$ in arbitrary order. The preference lists of the vertices in \mathcal{I}' are shown in Figure 2.

u_i :	u_i 's list in O	y_1, y_2, \dots, y_{K+1}	rest of women in W	$(1 \leq i \leq n/2)$
w_j :	w_j 's list in O	x_1, x_2, \dots, x_{K+1}	rest of men in U	$(1 \leq j \leq n/2)$
x_i :	$w_1, w_2, \dots, w_{n/2}$	y_i	women in $Y \setminus \{y_i\}$	$(1 \leq i \leq K + 1)$
y_j :	$u_1, u_2, \dots, u_{n/2}$	x_j	men in $X \setminus \{x_j\}$	$(1 \leq j \leq K + 1)$

Figure 2: Preference lists in the constructed instance of MIN BP SM FORBIDDEN DEC.

Having described G' and O' completely, all that remains is to specify the set of forbidden edges P . Each man u_i has $K + 1$ forbidden edges adjacent to him, namely, all edges to the newly-introduced y_1, y_2, \dots, y_{K+1} vertices. Similarly, edges between every w_j and all x_1, x_2, \dots, x_{K+1} vertices are also forbidden. In total, \mathcal{I}' has $n(K + 1)$ forbidden edges.

Claim 3.3. *If M is a perfect matching in \mathcal{I} admitting at most K blocking edges, then there is a matching M' in \mathcal{I}' with $M' \cap P = \emptyset$ admitting also at most K blocking edges.*

The construction of M' begins with copying M to G' . Since M is a perfect matching, all vertices in $V(G)$ are matched to vertices in $V(G)$ and thus, no forbidden edge can be in M' . The remaining vertices x_1, x_2, \dots, x_{K+1} and y_1, y_2, \dots, y_{K+1} are paired with each other: each $x_i y_i$ is added to M' .

M' is a perfect matching in G' , not containing any of the forbidden edges. Next, we show that no edge in $E(G') \setminus M'$ blocks M' that did not block M already. First of all, the forbidden edges do not block M' , because the preference lists of the vertices already in $V(G)$ were constructed in such a way that the vertices on preference lists in O are better than the vertices in $X \cup Y$, and all u_i, w_j vertices were matched in the perfect matching M . The first $n/2$ choices of any newly-added vertex in $X \cup Y$ are thus not blocking edges. At the same time, all these new vertices are matched to their first-choice partners among the newly-added vertices. Therefore no edge incident to them can block M' . All that remains is to observe that $u_i w_j$ edges blocking M' in \mathcal{I}' already blocked M in \mathcal{I} , because M is the restriction of M' to G . Therefore, the edges blocking M and M' are identical.

Claim 3.4. *If M' is a matching in \mathcal{I}' with $M' \cap P = \emptyset$ admitting at most K blocking edges, then its restriction to G is a perfect matching M in \mathcal{I} admitting at most K blocking edges.*

First, we discuss some essential structural properties of M' . The forbidden edges are not in M' , and at most K of them can block it. Suppose that there is a man u_i not married to any woman w_j in matching M . Since w_j ranks exactly $K + 1$ forbidden edges after its listed partners in G , and forbidden edges are the first $n/2$ choices of their other end vertex, all $K + 1$ of them block M' , regardless of the remaining edges in M' . Having derived a contradiction to our assumption that at most K edges block M' in total, we can state that each man u_i is matched in M' to a vertex w_j in O . Thus, the restriction of M' to G is a perfect matching with at most K blocking edges.

NP-hardness can be obtained for MIN BP SM FORCED DEC by simply forcing all edges of the form $x_i y_i$ in the above reduction. \square

We now strengthen Theorem 3.1 by giving strong lower bounds for the approximability of MIN BP SM FORBIDDEN and MIN BP SM FORCED. The reduction given in the proof of the next theorem builds on the reduction given in the proof of Theorem 3.1.

Theorem 3.5. *Each of MIN BP SM FORBIDDEN and MIN BP SM FORCED is not approximable within a factor of $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless $P = NP$. The result holds even if all preference lists are complete.*

Proof. We will give a reduction from the following NP-complete problem:

Problem 3.6. EXACT MAXIMAL MATCHING

Input: $\mathcal{I} = (G, K)$ comprising a bipartite graph G and a positive integer K .

Question: Is there a maximal matching M in G such that $|M| = K$?

EXACT MAXIMAL MATCHING is NP-complete even for graphs where all vertices representing men have degree two, while all vertices of the other colour class have degree three [26]. We show that if there were a polynomial approximation algorithm within a factor of $n^{1-\varepsilon}$ to MIN BP SM FORBIDDEN, then it would also find an exact maximal matching in \mathcal{I} .

In our proof, every instance $\mathcal{I} = (G, K)$ of EXACT MAXIMAL MATCHING is transformed into an instance $\mathcal{I}'' = (G', O', P)$ of MIN BP SM FORBIDDEN. We later show how to adapt the proof for MIN BP SM FORCED. Let n_1 and n_2 denote the size of each colour class in \mathcal{I} , such that $m = 2n_1 = 3n_2$.

We show that if there were a polynomial approximation algorithm with a performance guarantee of $n^{1-\varepsilon}$ for MIN BP SM FORBIDDEN (where n is the number of vertices in G'), then it would solve EXACT MAXIMAL MATCHING in polynomial time. To do so, another transformation is used, involving \mathcal{I}' , an instance of MIN BP PSMI. In [6], an instance $\mathcal{I}' = (G, O)$ of MIN BP PSMI is created from \mathcal{I} with special properties. One of them is that if G has a maximal matching of cardinality K , then \mathcal{I}' has a perfect matching admitting exactly $n_1 + n_2$ blocking edges. Otherwise, if G has no maximal matching of cardinality K , then any perfect matching in \mathcal{I}' is blocked by at least $n_1 + n_2 + C$ edges, where C is a huge number. To be more precise, let $B = \left\lceil \frac{3}{\varepsilon} \right\rceil$ and $C = (n_1 + n_2)^{B+1} + 1$ (in [6], the value of B was the same to that used here, but the value of C was slightly different). The number of vertices in each colour class of \mathcal{I}' is $3n_1 + 2mC + 4n_2 - K$.

Now we describe how \mathcal{I}' is transformed into \mathcal{I}'' . Note that this method is very similar to the one we used in the proof of Theorem 3.1. Denote by U and W the set of men and women in \mathcal{I}' , and let u_i and w_j denote an arbitrary man and woman in \mathcal{I}' respectively. Add these vertices to \mathcal{I}'' and then introduce C new men, namely $X = \{x_1, x_2, \dots, x_C\}$, and C new women, namely $Y = \{y_1, y_2, \dots, y_C\}$. Then each colour class in \mathcal{I}'' consists of $3n_1 + 2mC + 4n_2 - K + C$ vertices. The preference lists of the vertices in \mathcal{I}'' are shown in Figure 3.

u_i :	u_i 's list in O	y_1, y_2, \dots, y_C	rest of women in W	$(1 \leq i \leq U)$
w_j :	w_j 's list in O	x_1, x_2, \dots, x_C	rest of men in U	$(1 \leq j \leq W)$
x_i :	all women in W	y_i	women in $Y \setminus \{y_i\}$	$(1 \leq i \leq C)$
y_j :	all men in U	x_j	men in $X \setminus \{x_j\}$	$(1 \leq j \leq C)$

Figure 3: Preference lists in the constructed instance of MIN BP SM RESTRICTED.

The set of forbidden edges comprises all edges of the form $u_i x_j$ or $w_j y_i$. For MIN BP SM FORCED, the set of forced edges consists of all edges of the form $x_i y_i$. Due to this construction, and as in the proof of Theorem 3.1, if M is a matching in \mathcal{I}'' in which there is a man u_i not matched to a woman in O , then M is blocked by at least C edges.

It follows that if G has a maximal matching of size at most K then \mathcal{I}'' has a matching with at most $n_1 + n_2$ blocking pairs. On the other hand if G has no maximal matching of size at most K then any matching in \mathcal{I}'' has at least $C > (n_1 + n_2)^{B+1}$ blocking pairs. Hence an $(n_1 + n_2)^B$ -approximation algorithm for MIN BP SM RESTRICTED or MIN BP SM FORCED could be used to solve EXACT MAXIMAL MATCHING in polynomial time.

To complete the proof it remains to show that if n is the number of vertices in \mathcal{I}'' , then $n^{1-\varepsilon} \leq (n_1 + n_2)^B$. Using Inequalities 1-7 we give an upper bound for n , whilst with Inequalities 8-10 we establish a lower bound. Then, combining these two in Inequalities 11-14, we derive that $n^{1-\varepsilon} \leq (n_1 + n_2)^B$. Explanations for the steps are given as necessary after each of the three sets

of inequalities.

$$n = 2(3n_1 + 2mC + 4n_2 - K + C) \quad (1)$$

$$= 6n_1 + 8n_1C + 8n_2 - 2K + 2C \quad (2)$$

$$\leq 6n_1 + 8n_1((n_1 + n_2)^{B+1} + 1) + 8n_2 - 2K + 2(n_1 + n_2)^{B+1} + 2 \quad (3)$$

$$\leq 14n_1 + (n_1 + n_2)^{B+1}(8n_1 + 2) + 8n_2 + 2 \quad (4)$$

$$\leq 14n_1 + 14n_2 + (n_1 + n_2)^{B+1}(14n_1 + 14n_2) \quad (5)$$

$$\leq (n_1 + n_2)^{B+2} + 14(n_1 + n_2)^{B+2} \quad (6)$$

$$= 15(n_1 + n_2)^{B+2} \quad (7)$$

In (2) we use that $m = 2n_1$, whilst in (3) we use that $C = (n_1 + n_2)^{B+1} + 1$ by definition. To obtain (4) we omit $-2K$ whilst in (5) we assume that $n_2 \geq 1$ and increase all coefficients to the highest coefficient of 14. In (6) we assume that $n_1 \geq 1$, since $B \geq 3$.

$$n = 6n_1 + 8n_1C + 9n_2 - 2K + 2C \quad (8)$$

$$> (n_1 + n_2)^{B+1} \quad (9)$$

$$\geq 15^B \quad (10)$$

In (8) we keep only C from the right-hand side of the equality above and use the fact that $C > (n_1 + n_2)^{B+1}$, whilst in (10) we assume without loss of generality that $n_2 \geq 6$ so $n_1 \geq 9$ (recall that $2n_1 = 3n_2$).

$$C > (n_1 + n_2)^B \quad (11)$$

$$\geq 15^{-\frac{B}{B+2}} n^{\frac{B}{B+2}} \quad (12)$$

$$\geq n^{1-\frac{3}{B+2}} \quad (13)$$

$$\geq n^{1-\varepsilon} \quad (14)$$

Here (12) follows by (1)-(7); (13) follows by (8)-(10) and (14) uses the fact that $B \geq \frac{3}{\varepsilon}$. \square

3.2. Bounded parameters

Our results presented so far show that MIN BP SM RESTRICTED is computationally hard even if $P = \emptyset$ or $Q = \emptyset$. Yet if certain parameters of the instance or the solution can be considered as a constant, the problem can be solved in polynomial time. Theorem 3.7 firstly shows that this is true for MIN BP SM FORBIDDEN.

Theorem 3.7. MIN BP SM FORBIDDEN is solvable in $O(m^{L+1})$ time, where $L = |P|$, which is polynomial if L is a constant.

Proof. We firstly show how to solve MIN BP SM FORBIDDEN DEC in polynomial time. We assume that, for the purposes of this proof, the problem definition is modified so that, given an instance $\mathcal{I} = (G, O, P, K)$, we are required to find a matching M in G such that $M \cap P = \emptyset$ and $|bp(M)| \leq K$, or report that no such matching exists.

Our first observation is that this problem is trivially solvable if the target value K satisfies $K \geq L$. In this case, deleting the L forbidden edges from $E(G)$ and finding a stable matching in the remaining graph delivers a matching that is blocked in the original instance by only a subset of the removed edges (if any). Thus, a matching M with $M \cap P = \emptyset$ and $|bp(M)| \leq L \leq K$ always exists.

Now assume that $K < L$. Suppose firstly that there is a matching M with $M \cap P = \emptyset$ and $|bp(M)| = k \leq K < L$. If those k blocking edges are deleted from $E(G)$, then M is a stable matching in the remainder of G , and M contains none of the forbidden edges. Note that we did not specify which edges block M : they can be both forbidden and unrestricted.

Hence to solve MIN BP SM FORBIDDEN DEC we generate all subsets S of potential blocking edges, where $|S| \leq K$. After deleting the edges in S from G , we try to find, in the remaining graph, a matching M such that $M \cap P = \emptyset$ and M is stable, or we report that no such matching exists. By Theorem 1.2, this step can be accomplished in $O(m)$ time. If such a matching M exists, then it admits at most K blocking edges in \mathcal{I} .

Thus $\sum_{i=0}^K \binom{m}{i} = \sum_{i=0}^L \binom{m}{i}$ subsets are generated to determine whether the desired matching exists. The number of rounds is thus $O(m^L)$, while each round takes $O(m)$ time to complete. The overall running time is $O(m^{L+1})$.

We now show how to use the above approach in order to solve MIN BP SM FORBIDDEN. If we find a solution during the course of this process then G admits a matching M such that $M \cap P = \emptyset$ and $|bp(M)| \leq L$. In order to minimise $|bp(M)|$ it suffices to use the technique in the previous paragraph in combination with a binary search procedure on values of $K \leq L$. This requires $O(\log L)$ invocations of the algorithm for the decision problem, which is a constant, and hence the overall time complexity remains $O(m^{L+1})$. \square

In sharp contrast to the previous result on polynomial solvability when the number of forbidden edges is small, we state the following theorem for MIN BP SM FORCED DEC.

Theorem 3.8. MIN BP SM FORCED DEC is NP-complete even if $|Q| = 1$.

Proof. The NP-complete problem we reduce to MIN BP SM FORCED DEC is EXACT MAXIMAL MATCHING. As previously mentioned, this problem is NP-complete even for graphs where all vertices representing men have degree two, while all vertices of the other side have degree three [26]. Hence suppose we are given an instance $\mathcal{I} = (G, K)$ of this restriction, where in G , U_0 and W_0 are the two colour classes and E is the edge set.

In this proof, we construct a MIN BP SM FORCED DEC instance $\mathcal{I}' = (G', O', Q, K')$ with a single forced edge in such a way that there is a maximal matching of cardinality K in \mathcal{I} if and only if there is a matching containing the forced edge and admitting exactly $K' = |U_0| + |W_0|$ blocking edges in \mathcal{I}' . Our construction is based on ideas presented in [6].

All vertices in G rank their edges in an arbitrary but fixed order. We will refer to these labels when constructing \mathcal{I}' . We now describe \mathcal{I}' . The vertex set of graph G' in \mathcal{I}' can be partitioned into seven sets: U, V, W, Z, S_1, S_2, X and Y , where $U \cup V \cup X \cup S_1$ are the men and $W \cup Z \cup Y \cup S_2$ are the women. Specific subgraphs of G' are referred to as u -gadgets, w -gadgets, together with a special gadget containing the forced edge; see Figure 4. Aside from these, G' also contains some extra vertices, the so-called *garbage collectors*, partitioned into two sets: X and Y . Later we will see that these garbage collectors are paired to the vertices not covered by the matching in G . To that end, $|X| = |W_0| - K$ and $|Y| = |U_0| - K$. The whole construction is illustrated in Figure 5.




Figure 4: A u -gadget, a w -gadget and the special gadget

Each u -gadget replaces a vertex $u \in U_0$ in G . It is defined on five vertices: $u_1, u_2, u_3 \in U$ and $z_1, z_2 \in Z$. Its edges and the preferences on them are shown in Figure 4. Two *interconnecting edges* connect the special gadget to u_3 , and an interconnecting edge connects the special gadget to each of u_1 and u_2 . These edges are ranked last in the case of u_1 and u_2 , and ranked as the last two edges by u_3 . It is described later which vertices of the special gadget are incident to these interconnecting edges. The u -gadget also has edges to all w -gadgets representing vertices in W_0 to which u was adjacent. After describing the w -gadget, we elaborate on the position of these edges, referred to as *relevant edges*. Aside from these, every u_1 has edges to all garbage collectors in Y . These edges are all worse than the relevant edges of u_1 and they are ranked arbitrarily at the bottom of u_1 's list. The vertices in Y also rank all u_1 vertices arbitrarily.

The w -gadgets are structured similarly. Each gadget consists of seven vertices: $w_1, w_2, w_3, w_4 \in W$ and $v_1, v_2, v_3 \in V$. Aside from the edges within the gadget, it has two interconnecting edges between w_4 and vertices in the special gadget (described in detail later), and three relevant edges between w_1, w_2, w_3 and vertices of u -gadgets. These are the edges drawn in accordance with the edge labels. Suppose in \mathcal{I} , edge uw was ranked i th by u and j th by w , where $i \in \{1, 2\}$ and $j \in \{1, 2, 3\}$. Then in \mathcal{I}' , u_i in the u -gadget is connected to w_j in the w -gadget. Therefore, each edge in \mathcal{I} is transformed into a single edge in \mathcal{I}' and each u_i , $i \in \{1, 2\}$ and w_j , $j \in \{1, 2, 3\}$, has exactly one relevant edge. All of these edges are second choices of both of their end vertices. In addition to these, if $uw \in E(G)$, but the corresponding u - and a w -gadgets are not yet connected by u_1w_1 , we add u_1w_1 , which is referred as an *adjacency edge*. This edge is ranked by both u_1 and w_1 after their relevant edges, but ahead of their edges to garbage collectors. Similar to u -gadgets, w -gadgets are also connected to garbage collectors. Each w_1 vertex has $|W_0| - K$ edges to the vertices in X , ranked arbitrarily at the bottom of w_1 's preference list. Also the vertices in X rank the w_1 vertices arbitrarily.

The special gadget is defined on only six vertices in the set $S_1 \cup S_2$, where $S_1 = \{u_0, u'_0, u''_0\}$ and $S_2 = \{w_0, w'_0, w''_0\}$. The unique forced edge in the entire instance is u_0w_0 . Apart from u'_0 and w''_0 , they are connected to u - and w -gadgets. In each u -gadget, u_3 is adjacent to w_0 and w'_0 , and each of u_1 and u_2 is adjacent to w_0 . In each w -gadget, w_4 is adjacent to u_0 and u''_0 , and each of w_1, w_2 and w_3 is adjacent to u_0 . Moreover, u_0 and w_0 are connected to all garbage collectors of the opposite side via additional interconnecting

edges. These edges are ranked last by the vertices in X and Y . The four vertices u_0, u''_0, w_0, w'_0 prefer their interconnecting edges to their edges inside of the special gadget.




Figure 5: As the dashed grey relevant edge u_2w_2 shows, u and w were connected in \mathcal{I} by an edge ranked second by both of them. The dotted grey edge u_1w_1 is an adjacency edge. Some interconnecting edges have been omitted to avoid clutter.

Claim 3.9. *Corresponding to each maximal matching M in \mathcal{I} of cardinality K there is a matching M' in \mathcal{I}' with $u_0w_0 \in M'$ and $|bp(M')| = |U_0| + |W_0|$.*

Proof. First, the set of relevant edges in G' corresponding to M is chosen.

They cover exactly K of the $|U| = 3|U_0|$ vertices of U , and analogously, exactly K of the $|W| = 4|W_0|$ vertices in W .

In u -gadgets, where either of u_1 and u_2 has a relevant edge in M' , the other vertex in U is matched to its copy in Z . The remaining two vertices of the gadget are then paired to each other. In the other case, if u was unmatched in M , then $\{u_2z_2, u_3z_1\} \subseteq M'$, and $M'(u_1) \in Y$. Given the set of u_1 vertices to pair with the garbage collectors in Y , we find any stable matching in this subgraph and add it to M' . Note that this step matches the $|U_0| - K$ u_1 vertices to the $|U_0| - K$ garbage collectors in Y .

The strategy is similar for the w -gadgets. Suppose that some w_j is already matched to a vertex in U , because that relevant edge corresponds to a matching edge in M . In M' we then match w_4 with v_j and pair the remaining two vertices in W with their partners in V . Otherwise, if w was unmatched in M , then in M' , w_1 is matched to a garbage collector, and $\{w_2v_2, w_3v_3, w_4v_1\} \subseteq M'$. In the subgraph induced by the garbage collectors in X and the w_1 vertices corresponding to unmatched w vertices we construct a stable matching and add it to M' . This step matches the $|W_0| - K$ w_1 vertices to the $|W_0| - K$ garbage collectors in X .

In the special gadget, $u_0w_0, u'_0w'_0$ and $u''_0w''_0$ are chosen.

Now we investigate the number of blocking edges incident to at least one vertex in any u -gadget. The edges running to garbage collectors cannot block, because M' restricted to that subgraph is a stable matching and u_1 vertices not matched to garbage collectors have better relevant edges in M' . Since all u_3 vertices are matched to their first or second choices, their edges to the special gadget do not block either. Consider now a relevant edge $u_iw_j \notin M'$. Since M was a maximal matching, either u or w is matched in M . By construction of M' , if u is matched in M then u_i prefers $M'(u_i)$ to w_j , whilst if w is matched in M then w_j prefers $M'(w_j)$ to u_i . Regarding the adjacency edges, they only block M' if both of their end vertices are matched to garbage collectors. But they both are then unmatched and adjacent in G , which contradicts to the fact that M is maximal. The only edges remaining are in the u -gadgets. In each u -gadget, exactly one edge blocks M' : if u was matched to its i th ranked edge in M , then u_iz_i blocks M' , otherwise u_1z_1 blocks M' . Therefore, up to this point, we have exactly $|U_0|$ blocking edges.

Analogous arguments prove that among the edges incident to vertices in all w -gadgets, $|W_0|$ are blocking. In the previous paragraph we discussed that no relevant or adjacency edge blocks M' . The subgraph induced by the garbage collectors and w_1 vertices does not contain any blocking edge,

because a stable matching was chosen and the w_1 vertices not matched to garbage collectors are all matched in M' to a better vertex. Edges connecting w_4 vertices and the special gadget are last choice edges of the matched w_4 vertices. In the w -gadget, exactly one edge blocks M' : if w was matched and therefore $u_i w_j \in M'$, then $w_j v_j$, otherwise $w_1 v_1$.

It is easy to see that in the special gadget, none of the four non-matching edges blocks M' . ■

Claim 3.10. *Corresponding to each matching M' in \mathcal{I}' with $u_0 w_0 \in M'$ and $|bp(M')| = |U_0| + |W_0|$ there is a maximal matching M in \mathcal{I} of cardinality K .*

Proof. First we show that if $u_0 w_0 \in M'$, then each u - and w -gadget is adjacent to at least one blocking edge. Since w_0 prefers all its edges to $u_0 w_0$, if u_3 is not matched in M' to its first or second choice edge then $u_3 w_0$ blocks M' . But then if $u_3 z_i \in M'$ for $i \in \{1, 2\}$, it follows that $u_{3-i} z_{3-i}$ blocks M' . The same argument applies to u_0 and w_4 . If $M(w_4) \in V$, then $M(w_4)$ has a blocking edge, otherwise $u_0 w_4$ blocks M' . Therefore, if $|bp(M')| \leq |U_0| + |W_0|$, then each u - and w -gadget is incident to exactly one blocking edge.

If $u_3 w'_0 \in M'$ for some u -gadget then each of $u_3 w_0$ blocks and $u'_0 w_0$ blocks M' . Thus $|bp(M')| \geq |U_0| + |W_0| + 1$, a contradiction. Thus $u_3 w'_0 \notin M'$ for any u -gadget. By a similar argument we can establish that $u''_0 w_4 \notin M'$ for any w -gadget. If u_3 is unmatched in M' for some u -gadget then each of $u_3 w_0$ and $u_3 w'_0$ blocks M' , since $u'_3 w'_0 \notin M'$ for any u' -gadget. Again that would imply that $|bp(M')| \geq |U_0| + |W_0| + 1$, a contradiction. By a similar argument, w_4 is matched in M' for each w -gadget.

We have established that, for each u -gadget, exactly one of $u_i z_i$ blocks M' for some $i \in \{1, 2\}$, and for each w gadget, exactly one of $v_j w_j$ blocks M' for some $j \in \{1, 2, 3\}$, and these are the only blocking pairs of M' in \mathcal{I}' . Hence in each u -gadget, u_i is matched in M' for $i \in \{1, 2\}$, for otherwise $u_i w_0$ blocks M' . Similarly in each w -gadget, w_j is matched in M' for $j \in \{1, 2, 3\}$. It also follows by a similar argument that each member of $X \cup Y$ is matched in M' . Finally we can observe that $u'_0 w'_0 \in M'$ for otherwise $u'_0 w_0$ blocks M' , and $u''_0 w''_0 \in M'$ for otherwise $u_0 w''_0$ blocks M' .

Thus in each u -gadget, $M'(u_3) = z_i$ for some $i \in \{1, 2\}$ whilst $M'(w_4) = v_j$ for some $j \in \{1, 2, 3\}$. It follows that in M' , u_i is matched either via a relevant edge or adjacency edge, or to a garbage collector in Y . Similarly in M' , w_j is matched either via a relevant edge or adjacency edge, or to a garbage collector in X . Meanwhile u_{3-i} is matched in M' to his first-choice vertex in Z , whilst the two vertices in $\{w_1, w_2, w_3\} \setminus \{w_j\}$ are matched in M' to their

first-choice vertices in V . Define a set of edges M in G as follows:

$$M = \{uw \in E : u_i w_j \in M' \text{ for some } i \in \{1, 2\} \text{ and } j \in \{1, 2, 3\}\}.$$

Then M is a matching in G and moreover $|M| = K$, since all $|U_0| - K$ vertices in Y are matched in M' to u_1 vertices of various u -gadgets, and similarly, all $|W_0| - K$ vertices in X are matched in M' to w_1 vertices. That leaves K of the u -gadgets that contribute a single relevant edge (or adjacency edge) to M' . All that remains to show is that this matching is maximal. Let us suppose otherwise, i.e., there are two gadgets corresponding to vertices u and w in G such that all their vertices in U and W in G' are matched to either garbage collectors or to their z - or v -copies. This is only possible if u_1 and w_1 are both matched to garbage collectors, but then the adjacency edge $u_1 w_1$ blocks M' . \square

A counterpart to Theorem 3.7 holds in the case of MIN BP SM RESTRICTED if the number of blocking pairs in an optimal solution is a constant.

Theorem 3.11. *MIN BP SM RESTRICTED is solvable in $O(m^{L+1})$ time, where L is the minimum number of edges blocking an optimal solution, which is polynomial if L is a constant.*

Proof. Analogously to the proof of Theorem 3.7, we show how to solve MIN BP SM RESTRICTED DEC in polynomial time. Again, we assume for the purposes of this proof that the problem definition is modified so that, given an instance $\mathcal{I} = (G, O, P, Q, K)$, we are required to find a matching M in G such that $M \cap P = \emptyset$, $Q \subseteq M$ and $|bp(M)| \leq K$, or report that no such matching exists.

Suppose there exists a matching M in G such that $M \cap P = \emptyset$, $Q \subseteq M$ and $|bp(M)| \leq K$. If we form G' by deleting the edges in $bp(M)$ then M is stable in G' . Hence to find M it suffices to generate all subsets S of edges of size at most K and form a graph G_S by deleting S from G . We can determine in linear time whether G_S admits a stable matching M' that satisfies all constraints on restricted edges [7]. If so, $|bp(M')| \leq K$ in G . There are $O(Km^K)$ sets of edges to remove and checking the existence of a stable matching satisfying constraints on restricted edges can be done in $O(m)$ time.

We now show how to use the above approach to solve MIN BP SM RESTRICTED. For each value of K , where K starts from 0 and increases by

1 after each iteration, we execute the algorithm in the previous paragraph, noting that it is sufficient to generate all subsets of size exactly K at each iteration. We terminate as soon as we find a matching M satisfying the constraints on restricted edges such that $|bp(M)| \leq K$. This process is bound to halt, since by definition, \mathcal{I} admits a matching M such that $M \cap P = \emptyset$, $Q \subseteq M$ and $|bp(M)| = L$, so $K \leq L$. Thus the overall time complexity of this approach is $O((L+1)m^{L+1}) = O(m^{L+1})$, which is polynomial if L is a constant. \square

Next we study the case of degree-constrained graphs; for most hard SM and SR problems, it is the most common special case to investigate [21, 17, 5]. Here, we show in Theorem 3.12 that MIN BP SM RESTRICTED remains computationally hard even for instances with preference lists of length at most 3. On the other hand, according to Theorem 3.17, the problem can be solved in polynomial time when the length of preference lists is bounded by 2.

Theorem 3.12. *MIN BP SM FORBIDDEN DEC and MIN BP SM FORCED DEC are NP-complete. The result holds even if each agent's preference list is of length at most 3.*

Proof. We give a reduction from (2,2)-E3-SAT, a restriction of SATISFIABILITY, which may be defined as follows:

Problem 3.13. (2,2)-E3-SAT

Input: a Boolean formula B in CNF, in which each clause comprises exactly 3 literals and each variable appears exactly twice unnegated and exactly twice negated.

Question: Is there a satisfying truth assignment for B ?

(2,2)-E3-SAT is NP-complete [3]. Given an instance B of this problem, let us denote the number of variables by n_B and the number of clauses by m_B .

Using the simple transformation described in Section 2, any MIN BP SM FORCED instance can be converted into a MIN BP SM FORBIDDEN instance without increasing the preference list lengths. Hence it is sufficient to investigate MIN BP SM FORBIDDEN.

Our goal is to construct an instance $\mathcal{I} = (G, O, P)$ of MIN BP SM FORBIDDEN such that B is satisfiable if and only if \mathcal{I} admits a matching M with $M \cap P = \emptyset$ and $|bp(M)| \leq n_B + m_B$.

Our construction combines ideas from two papers. Corresponding to B , we introduce a variable gadget and a clause gadget. The first one is a slightly more sophisticated variant of the variable gadget used in Theorem 7 of [6], to show NP-hardness of finding a maximum cardinality almost stable matching. Our clause gadget is a simplified version of another clause gadget from Theorem 1 in [5]. There, the Almost Stable Roommates problem is shown to be NP-hard. Both proofs investigate the case with bounded preference lists.

When constructing instance \mathcal{I} from the given Boolean formula B , we keep track of the order of the three literals in each clause and the order of the two unnegated and two negated appearances of each variable.




Figure 6: A clause and a variable gadget with their special matchings, marked by grey edges. The dotted edges are forbidden.

The variable gadget. For each variable in B , a *variable gadget*, which is a graph on 44 vertices, is defined. The right hand-side of Figure 6 illustrates the essential part of such a gadget, which is a cycle of length 24. This cycle contains no forbidden edges, and each vertex along it has degree 3 due to being incident to an additional forbidden edge. For sake of simplicity, only four of these forbidden edges are depicted in the figure, namely the ones incident to vertices x_1, x_2, x_3 and x_4 , as they are responsible for the communication between clause and variable gadgets. Each of these four vertices has its third, forbidden edge connected to a clause gadget. These edges are called *interconnecting edges* and are ranked second on the preference list of both of their end vertices.

Consider a variable v in B . Due to the properties of $(2,2)$ -E3-SAT, v occurs twice in unnegated form, say, in clauses C_1 and C_2 . Suppose that v 's first unnegated appearance is as the i th literal of C_1 . This is represented by the interconnecting edge between x_1 (in the vertex gadget corresponding to v) and vertex a_i in the clause gadget (described below) corresponding to C_1 . Similarly, x_2 is connected to an a -vertex in the clause gadget of C_2 . The same variable, v , also appears twice in negated form. The relevant variables in the gadgets representing those clauses are connected to x_3 and x_4 . The other end vertices of these two interconnecting edges mark where these two literals appear in their clauses.

As mentioned before, all vertices along the cycle have exactly one forbidden edge attached to them. In the case of x_1, x_2, x_3 and x_4 , these edges take the form of interconnecting edges. Regarding the remaining 20 vertices of the cycle, there is a dummy vertex with a forbidden edge attached to each of them, which we call *variable pendant edges* (these edges are not depicted in Figure 6). This edge is their last choice. These edges guarantee that if any of these 20 vertices remain unmatched in the cycle, then there is a contribution of one blocking edge.

Two special matchings are defined on a variable gadget: M_T , denoted by grey edges and M_F , comprising the black edges. While M_T is blocked by x_4u_8 exclusively, M_F is blocked by x_1u_1 exclusively.

Claim 3.14. *Let M be a matching on a variable gadget. If M is not M_T or M_F , then it is blocked by at least two edges belonging to the variable gadget.*

Proof. Since x_1u_1 and x_4u_8 are best-choice edges of both of their end vertices,

they block any matching not containing them. If both of them are in M , then there is at least one unmatched vertex on the path between u_8 and u_1 via y_4 , and another unmatched vertex on the path between x_4 and x_1 via y_2 . The first path comprises vertices with a forbidden edge as their last choice, therefore, it contributes a blocking edge. The only way to avoid additional blocking edges on the second path is to leave either x_2 or x_3 unmatched. Since they both are first choices of some other vertex along the cycle, we obtain a second blocking edge.

The remaining case is when exactly one of x_1u_1 and x_4u_8 is in M . If every second edge in the cycle belongs to M , then it is either M_T or M_F . Otherwise, for simple parity reasons, there are two unmatched vertices on the 24-cycle. Suppose firstly that $x_1u_1 \in M$. Then $x_4u_8 \notin M$ so x_4u_8 blocks M . If a u -, v - or y -vertex is unmatched in M then we obtain a further blocking edge from the vertex's variable pendant edge. Otherwise either x_2 is unmatched, so x_2u_4 blocks M , or x_3 is unmatched, so x_3u_5 blocks M , or x_4 is unmatched, so x_4u_7 blocks M . The argument is similar if $x_4u_8 \in M$. ■

The clause gadget. To each clause in B , a graph on 14 vertices is defined, which we refer to as the *clause gadget*. Three of them, a_1, a_2 and a_3 , are connected to variable gadgets via interconnecting edges, all ranked second. There are three special matchings on a clause gadget, blocked by only a single edge. They can be seen on the left hand-side of Figure 6. The grey edges selected in each of the clause gadget copies from top to bottom denote M_1, M_2 and M_3 respectively. In addition to the interconnecting edges, a clause gadget has two forbidden edges, namely p_3r_1 and q_3r_2 .

Claim 3.15. *Let M be a matching in a clause gadget. If M is not M_1 , M_2 or M_3 , then it is blocked by at least two edges, both of them belonging to the clause gadget.*

Proof. First, suppose that $M \neq M_i$ for all $i \in \{1, 2, 3\}$ and that M is blocked by at most one edge. Since all three edges connecting a and b -vertices are first choices of both of their end vertices, they block any matching not including them. Another restriction arises from the fact that the forbidden edges r_1p_3 and r_2q_3 ensure that if p_3 or q_3 is unmatched, they also contribute a blocking edge. Similarly, if p_1 or q_1 is unmatched, they contribute a blocking edge.

Suppose $b_ia_i \in M$ for all $i \in \{1, 2, 3\}$. Then, p_2 is matched either to p_1 or to p_3 , leaving the other one unmatched. The same argument applies for

the q -vertices on the other side of the gadget. Therefore, at least two edges from the clause gadget block M .

In the remaining case, exactly one of the $b_i a_i$ edges is outside of M . Since we are searching for a matching blocked by at most one edge, no further blocking edge can occur. Therefore, p_1 , q_1 , p_3 and q_3 are all matched in M . From this point on, it is easy to see that all matchings fulfilling these requirements are M_1 , M_2 and M_3 . ■

Claims 3.14 and 3.15 guarantee that if a matching M 's restriction to any of the variable or clause gadgets deviates from their special matchings, then $|bp(M)| > n_B + m_B$.

Claim 3.16. *B is satisfiable if and only if \mathcal{I} admits a matching M such that $M \cap P = \emptyset$ and $|bp(M)| \leq n_B + m_B$.*

Proof. Suppose we are given a satisfying truth assignment f for B . We construct a matching M in \mathcal{I} as follows. In the variable gadgets, the edges of M_T are chosen if the corresponding variable is **true** under f , and the edges of M_F are chosen otherwise. There is at least one literal in each clause that is **true** under f . If this literal is the i th in the clause, matching M_i is chosen, where $i \in \{1, 2, 3\}$. If more than one literal is true, we choose one of them arbitrarily. Clearly M contains no forbidden edges. It is not difficult to verify that each gadget contributes a single blocking edge. Thus $|bp(M)| \geq n_B + m_B$. As a last step, we show that no interconnecting edge blocks M , and thus $|bp(M)| = n_B + m_B$. Suppose that $a_i x_j$ blocks M . Since it is the second choice of a_i , it follows that $b_i a_i \notin M$. We now know that the i th literal of the clause was **true** in the truth assignment. Therefore, x_i is matched to its first choice.

To prove the converse direction, we utilise Claims 3.14 and 3.15. On one hand, these two statements, together with the characteristics of the special matchings, prove that $bp(\mathcal{I}) \geq n_B + m_B$. On the other hand, $|bp(M)| = n_B + m_B$ occurs if and only if M 's restriction to variable gadgets is M_T or M_F , and its restriction to clause gadgets is M_1 , M_2 or M_3 . Then, assigning **true** to all variables with M_T in their gadgets and **false** to the rest results in a truth assignment f . Since no interconnecting edge blocks M , at least one literal per clause is **true**, so f satisfies B . ■

Theorem 3.17. *MIN BP SM RESTRICTED is solvable in $O(n)$ time if each preference list consists of at most 2 elements.*

Proof. In this constructive proof we describe an algorithm that produces an optimal matching. First, the input is simplified. Then, the graph is segmented so that each subgraph falls into a category with a specified rule for selecting the edges of an optimal matching. As in previous cases, it is sufficient to tackle MIN BP SM FORBIDDEN, because instances of MIN BP SM FORCED can be transformed to this problem.

Due to the degree constraints, every component of the underlying graph is a path or a cycle. If any of these components is free of forbidden edges, then we simply fix a stable matching on it. This step is executed whenever such a component appears during the course of the algorithm. For those components with forbidden edges, we split all vertices having a first-choice forbidden edge and a second-choice edge – unrestricted or forbidden – into two vertices. This change does not affect $|bp(M)|$, because in this case, each edge blocks the matching if and only if its other end vertex is matched to a worse partner or is unmatched. After this splitting is executed, all components contain forbidden edges that start paths or that are inside a path, being the last choices of both of their end vertices.

Each component consists of *segments* of unrestricted edges, separated by forbidden edges. When talking about a segment, we always mean a series of adjacent unrestricted edges. Since unrestricted cycles have already been eliminated by fixing a stable matching on them, every segment is a path. Due to Theorem 1.1, each path admits a unique stable matching. Fixing a matching on a segment induces blocking edges only among the unrestricted edges of the segment and the forbidden edges adjacent to the segment. We claim that in an optimal solution, each segment and the (at most two) forbidden edges surrounding it contribute at most two blocking edges. This is simply due to the fact that any stable solution on the unrestricted edges is blocked only by forbidden edges. Therefore, deviating from this solution might only pay off if the matching restricted to this segment is blocked by a single edge and covers both of its end vertices.

The unique stable matching M on a segment $\langle v_1, v_2, \dots, v_k \rangle$ falls into exactly one of the following categories:

1. M covers both v_1 and v_k ;
2. M covers either v_1 or v_k ;
3. M covers neither v_1 nor v_k .

In each step of our algorithm, a segment is chosen and a matching is fixed on it. The segment and some of the forbidden edges adjacent to it are then removed from the graph. This is done in the following way in these three cases.

In case 1, an optimal solution arises from choosing M . If a forbidden edge e is incident to either v_1 or v_k , it cannot block M . Nor can it block any superset of M in the original instance, so e can be deleted. In case 2, again the optimal solution arises from choosing M . Without loss of generality suppose that v_1 is covered. As in case 1, if a forbidden edge is incident to v_1 , it cannot block a superset of M in the original instance. Now suppose that a forbidden edge e is incident to v_k . Edge e may block M , and may also block a superset of M in the original instance, so it is retained.

The third case is divided into two subcases, depending on whether there is a matching M' that is blocked by only one edge and covers both v_1 and v_k . Finding such a matching or proving that none exists can be done iteratively, assuming that a chosen edge is the single blocking edge and then constructing M' so that no more edge blocks it. If such an M' does not exist, then M is chosen, and the segment (but not the forbidden edges) is removed. In the end the matching restricted to this segment will be blocked by no edge other than the two forbidden edges. Suppose M is not an optimal choice in this segment. Then the optimal matching M'' has at most one blocking edge from this segment and its adjacent edges. M'' must cover at least one endpoint vertex of the segment (otherwise we have at least two blocking edges) – without loss of generality let that vertex be v_1 . M'' cannot cover v_k given the non-existence of the aforementioned matching M' that was sought. Hence the forbidden edge incident to v_k contributes a blocking edge. This implies that M'' restricted to the segment itself is a stable matching, contradicting the fact that the unique stable matching for this segment covers neither v_1 nor v_k .

The only remaining case is that, in the segment $\langle v_1, v_2, \dots, v_k \rangle$, the unique stable matching M covers neither v_1 nor v_k , but there is a matching M' blocked by one (unrestricted) edge that covers both v_1 and v_k , and this is true for all remaining segments. Let C be any remaining component, which is a path comprising segments S_1, S_2, \dots, S_r , where S_i and S_{i+1} are separated by a forbidden edge ($1 \leq i \leq r-1$), together with a possible additional forbidden edge at each end of the path that may be remaining. Let M_i denote a stable matching in S_i (which covers neither endpoint vertex, as previously noted), and let M'_i denote a matching in segment S_i that covers

both endpoint vertices and is blocked by at most one unrestricted edge. Now let M'' be an optimal matching in C . We will show how to transform M'' to M''' such that $|bp(M''')| \leq |bp(M'')|$, $M_i \cup M_{i+1} \not\subseteq M'''$ and $M'_i \cup M'_{i+1} \not\subseteq M'''$ for any i ($1 \leq i \leq r-1$). Firstly let $M''' = M''$. Iteratively from $i = 1$ to $r-1$ we modify M''' , if necessary, as follows. If $M_i \cup M_{i+1} \subseteq M''$ then replace M_i by M'_i in M''' . Alternatively if $M'_i \cup M'_{i+1} \subseteq M''$ then replace M'_{i+1} by M_{i+1} in M''' . It follows that M''' has the desired properties once this process terminates. Thus to reach an optimal solution in C it suffices to let M^1 be the union of M_1, M'_2, M_3, \dots , and let M^2 be the union of M'_1, M_2, M'_3, \dots , and pick whichever of the two admits the fewer blocking pairs in C . \square

Even with the previous two theorems, we have not quite drawn the line between tractable and hard cases in terms of vertex degrees. The complexity of MIN BP SM RESTRICTED remains open for the case when preference lists are of length at most 2 on one side of the bipartite graph and are of unbounded length on the other side. However we believe that this problem is solvable in polynomial time.

Conjecture 3.18. *MIN BP SM RESTRICTED is solvable in polynomial time if each woman's preference list consists of at most 2 elements.*

3.3. Stable Roommates problem

Having discussed several cases of SM, we turn our attention to non-bipartite instances. Since SM is a restriction of SR, all established results on the NP-hardness and inapproximability of MIN BP SM RESTRICTED carry over to the non-bipartite SR case. As a matter of fact, more is true, since MIN BP SR RESTRICTED is NP-hard and difficult to approximate even if $P = \emptyset$ and $Q = \emptyset$ [1]. We summarise these observations as follows.

Remark 3.19. *By Theorems 3.1 and 3.5, each of MIN BP SR FORBIDDEN and MIN BP SR FORCED is NP-hard and not approximable within $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless $P = NP$. Moreover Theorems 3.12 and 3.8 imply that each of MIN BP SR FORBIDDEN and MIN BP SR FORCED is NP-hard even if all preference lists are of length at most 3 or, in the latter case, $|Q| = 1$. Finally MIN BP SR RESTRICTED is NP-hard and not approximable within $n^{\frac{1}{2}-\varepsilon}$, for any $\varepsilon > 0$, unless $P = NP$, even if $P = \emptyset$ and $Q = \emptyset$ [1].*

Remark 3.19 already shows that Theorem 3.7 does not carry over to the SR case, since MIN BP SR FORBIDDEN is computationally hard even if $P = \emptyset$.

As for the other polynomially solvable cases, the proof of Theorem 3.11 carries over without applying any modifications. Theorem 3.17 also carries over to the SR case, but it needs a slight modification. If $\deg(v) \leq 2$ for every $v \in V(G)$, then G consists of paths and cycles. Each odd preference cycle without a forbidden edge contributes at least one blocking edge to any matching [32] and any maximal matching on such a cycle is blocked by exactly one edge. On the remainder of the graph, the algorithm described in the proof of Theorem 3.17 delivers an optimal matching for MIN BP SR RESTRICTED. The following remark summarises the discussed positive results.

Remark 3.20. *MIN BP SR RESTRICTED is solvable in polynomial time if the minimal number of edges blocking an optimal solution is a constant or if each preference list consists of at most 2 elements.*

4. Stable matchings with the minimum number of violated constraints on restricted edges

In this section, we study the second intuitive approximation concept. The desired matching is stable and violates as few constraints on restricted edges as possible. We return to our example that already appeared in Figure 1. As already mentioned earlier, the instance admits a single stable matching, namely $M = \{u_1w_1, u_2w_2, u_3w_3, u_4w_4\}$. Since M contains both forbidden edges, the minimum number of violated constraints on restricted edges is 2.

This section is structured as follows: in Section 4.1, complexity and approximability results are presented for SM MIN RESTRICTED VIOLATIONS, SR MIN FORBIDDEN, SR MAX FORCED and SR MIN RESTRICTED VIOLATIONS. In Section 4.2 we consider the complexity of SR MIN RESTRICTED VIOLATIONS when the degree of the underlying graph is bounded.

4.1. General complexity and approximability results

As mentioned in Section 1, a weighted stable matching instance models SM MIN RESTRICTED VIOLATIONS.

Theorem 4.1. *SM MIN RESTRICTED VIOLATIONS is solvable in polynomial time.*

Proof. We convert SM MIN RESTRICTED VIOLATIONS into a weighted SM problem using the following weight function:

$$w(e) = \begin{cases} -1 & \text{if } e \text{ is forced,} \\ 0 & \text{if } e \text{ is unrestricted,} \\ 1 & \text{if } e \text{ is forbidden.} \end{cases}$$

Then a stable matching M has weight $|M \cap P| - |M \cap Q| = |M \cap P| + |Q \setminus M| - |Q|$. Since $|Q|$ is invariant, a stable matching of minimum weight will violate the minimum number of constraints on restricted edges. The polynomial-time solvability of SM MIN RESTRICTED VIOLATIONS then follows by the fact that we can find a minimum weight stable matching in SM in polynomial time as discussed in the last paragraph of Section 2. \square

In the SR context, finding a minimum weight stable matching is NP-hard [8], so the above technique for SM does not carry over to SR. Indeed special cases of SR MIN RESTRICTED VIOLATIONS are NP-hard, as the following result shows.

Theorem 4.2. SR MIN FORBIDDEN DEC and SR MAX FORCED DEC are NP-complete.

Proof. Clearly both problems belong to the class NP. We begin by proving the NP-hardness of SR MIN FORBIDDEN DEC. To do so, we will use a reduction from the decision version of the following problem:

Problem 4.3. MIN VX COVER

Input: $\mathcal{I} = G$; a graph G on n vertices and m edges.

Output: A vertex cover $C \subseteq V(G)$ such that $|C| \leq |C'|$ for every vertex cover C' .

Specifically, define MIN VX COVER DEC to be the problem of deciding, given a graph G and an integer K , whether G admits a vertex cover of size at most K . MIN VX COVER DEC is NP-complete [14].

Given an instance (G, K) of MIN VX COVER DEC, the following instance (G', O, P, K) of SR MIN FORBIDDEN DEC is constructed. The entire graph G is copied, and then, a gadget is attached to each vertex $v_i \in V(G)$. It is a complete bipartite graph on four vertices: one of them is $p_i = v_i$, whilst the remaining three are denoted by \bar{p}_i, q_i and \bar{q}_i . Vertex p_i ranks \bar{p}_i at the top, followed by all p_j vertices such that $v_i v_j \in E(G)$, in arbitrary order, followed

by \bar{q}_i at the bottom of this list. The new vertices' orderings can be seen in Figure 7. In order to derive an instance with complete lists, all remaining vertices can be placed in arbitrary order to the bottom of the lists. Later we will see that these edges never appear in stable matchings, neither do they block them. The set of forbidden edges comprises all $p_i\bar{p}_i$ edges corresponding to the dotted grey edges in our illustrations in Figure 7.




Figure 7: Adding $K_{2,2}$ to each vertex of the MIN VX COVER DEC instance

Claim 4.4. *If M is a stable matching in G' , then for each i ($1 \leq i \leq n$), either $\{p_i\bar{p}_i, q_i\bar{q}_i\} \subseteq M$, or $\{p_i\bar{q}_i, \bar{p}_i q_i\} \subseteq M$.*

Proof. This claim follows from the structure of the introduced gadget. First, we observe that in M , each \bar{p}_i is either matched to p_i or to q_i , for otherwise $p_i\bar{p}_i$ blocks M , since \bar{p}_i is p_i 's first choice. Similarly, $q_i\bar{q}_i \in M$ or $q_i\bar{p}_i \in M$,

for otherwise $\bar{p}_i q_i$ would block M . Finally, $p_i \bar{q}_i \in M$ or $q_i \bar{q}_i \in M$, otherwise $q_i \bar{q}_i$ blocks M . These three requirements imply that the claim must hold. ■

Claim 4.5. *If there is a vertex cover $C \subseteq V(G)$ with $|C| \leq K$, then there is a stable matching M in G' for which $|M \cap P| \leq K$.*

Proof. The matching M is constructed based on the following case distinction:

$$\begin{aligned} \{p_i \bar{p}_i, q_i \bar{q}_i\} &\subseteq M & \text{if } v_i \in C \\ \{p_i \bar{q}_i, \bar{p}_i q_i\} &\subseteq M & \text{if } v_i \notin C \end{aligned}$$

Clearly $|M \cap P| \leq K$. Moreover, no edge in the gadgets can block M , because the preferences inside the gadget are cyclic. Due to the vertex cover property, edges between two p -vertices have at least one end vertex in C , thus, at least one of their end vertices is matched to its first-choice partner \bar{p} in G' . For each vertex in G' , the edges in the sets “rest” are worse than the edge in M . ■

Claim 4.6. *If there is a stable matching M in G' for which $|M \cap P| \leq K$, then there is a vertex cover $C \subseteq V(G)$ with $|C| \leq K$.*

Proof. Claim 4.4 allows us to investigate only two cases per gadget. We use a similar function to that in Claim 4.5, but in the opposite direction, in order to derive C from M , as follows:

$$\begin{aligned} v_i \in C & \text{ if } \{p_i \bar{p}_i, q_i \bar{q}_i\} \subseteq M \\ v_i \notin C & \text{ if } \{p_i \bar{q}_i, \bar{p}_i q_i\} \subseteq M \end{aligned}$$

Trivially, $|C| \leq K$. Suppose C is not a vertex cover. Then, there is an edge $v_i v_j = p_i p_j$ for which $\{p_i \bar{q}_i, p_j \bar{q}_j\} \subset \{p_i \bar{q}_i, \bar{p}_i q_i, p_j \bar{q}_j, \bar{p}_j q_j\} \subseteq M$. Then $p_i p_j$ blocks M . ■

To prove NP-completeness for SR MAX FORCED DEC, let (G', O, Q, K) be the constructed instance, where G' , O and K are defined as in the SR MIN FORBIDDEN DEC case, and $Q = \{p_i \bar{q}_i : 1 \leq i \leq n\}$. The correctness of the reduction then follows by a similar argument □

We remark that MIN VX COVER is NP-hard and cannot be approximated within a factor of $2 - \varepsilon$ for any $\varepsilon > 0$, assuming that the Unique Games Conjecture (UGC) is true [22]. When viewing the construction in Theorem 4.2 as a reduction from MIN VX COVER to SR MIN FORBIDDEN, the measures of an optimal solution in each problem instance are identical. The same is true in the case of SR MAX FORCED. It follows that each of SR MIN FORBIDDEN

and SR MAX FORCED is not approximable within a factor of $2 - \varepsilon$ for any $\varepsilon > 0$, assuming the UGC holds.

When studying SR MAX FORCED, we measured optimality by keeping track of the number of violated constraints. One might find it more intuitive instead to maximise $|Q \cap M|$, the number of forced edges in the stable matching. Our NP-hardness proof for SR MAX FORCED given in Theorem 4.2 can be used without modification to show NP-hardness under this different measure, however the approximability results need to be revisited. In fact, this modification of the measure changes the approximability of the problem as well:

Theorem 4.7. *For SR MAX FORCED, the maximum of $|Q \cap M|$ cannot be approximated within $n^{\frac{1}{2}-\varepsilon}$ for any $\varepsilon > 0$, unless $P = NP$.*

Proof. We adapt the proof of Theorem 4.2 so that the reduction is from MAX IND SET, the problem of finding a maximum independent set in a given graph $G = (V, E)$. MAX IND SET is not approximable within $N^{1-\varepsilon}$ for any $\varepsilon > 0$, unless $P = NP$ [35], where $N = |V|$. In the modified reduction an independent set of vertices S in G corresponds to the matching

$$M = \{p_i \bar{q}_i, \bar{p}_i q_i : v_i \in S\} \cup \{p_i \bar{p}_i, q_i \bar{q}_i : v_i \notin S\}$$

in the constructed instance \mathcal{I} of SR MAX FORCED. Suppose that A is an $n^{\frac{1}{2}-\varepsilon}$ -approximation algorithm that approximates $|Q \cap M_{opt}|$ in \mathcal{I} , for some $\varepsilon > 0$, where M_{opt} is an optimal solution and n is the number of agents in \mathcal{I} . Note that $|S_{opt}| = |M_{opt}|$, where S_{opt} is a maximum independent set in G . Moreover $n^{\frac{1}{2}-\varepsilon} = (4N)^{\frac{1}{2}-\varepsilon} \leq N^{1-\varepsilon}$ since $n = 4N$ and without loss of generality we can assume that $N \geq 4$. We thus reach a contradiction to the inapproximability of MAX IND SET. \square

4.2. Bounded parameters

We now turn to the complexity of SR MIN RESTRICTED VIOLATIONS and its variants when the degree of the underlying graph is bounded or some parameter of the instance can be considered as a constant. With Theorems 4.8 and 4.12 we draw the line between NP-hard and polynomially solvable cases in terms of degree constraints.

Theorem 4.8. *SR MIN FORBIDDEN DEC and SR MAX FORCED DEC are NP-complete even if every preference list is of length at most 3.*

Proof. As in Theorem 4.2, both problems belong to NP. We firstly show the NP-hardness of SR MIN FORBIDDEN DEC for this length restriction on preference lists. To do so, we reduce from MIN VX COVER DEC in cubic graphs, which is NP-complete [15, 25]. Hence let $\mathcal{I} = (G, K)$ be an instance of MIN VX COVER, where $G = (V, E)$ is a cubic graph, $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_m\}$. For each i ($1 \leq i \leq n$), suppose that v_i is incident to edges e_{j_1} , e_{j_2} and e_{j_3} in G , where without loss of generality $j_1 < j_2 < j_3$. Define $e_{i,r} = e_{j_r}$, where $r \in \{1, 2, 3\}$. Similarly for each j ($1 \leq j \leq m$), suppose that $e_j = v_{i_1}v_{i_2}$, where without loss of generality $i_1 < i_2$. Define $v_{j,s} = v_{i_s}$, where $s \in \{1, 2\}$.

We construct an instance \mathcal{I}' of SR MIN FORBIDDEN DEC as follows. The set $V' \cup W \cup E' \cup F$ constitutes the set of vertices in \mathcal{I}' , where these sets are defined as follows:

$$\begin{aligned} V' &= \{v_i^r : 1 \leq i \leq n \wedge r \in \{1, 2, 3\}\} \\ W &= \{w_i^r : 1 \leq i \leq n \wedge r \in \{1, 2, 3\}\} \\ E' &= \{e_j^s : 1 \leq j \leq m \wedge s \in \{1, 2\}\} \\ F &= \{f_j^s : 1 \leq j \leq m \wedge s \in \{1, 2\}\} \end{aligned}$$

The preference lists of these vertices (also indicating the edges of the graph) are as shown in Figure 8. In the preference list of a vertex v_i^r , the symbol $e(v_i^r)$ denotes vertex $e_j^s \in E'$ such that $e_j = e_{i,r}$ and $v_i = v_{j,s}$. Since i and r are given, e_j can be computed. Now we know i and j in the second equation, therefore we can compute s . Similarly in the preference list of vertex e_j^s , the symbol $v(e_j^s)$ denotes vertex $v_i^r \in V'$ such that $e_j = e_{i,r}$ and $v_i = v_{j,s}$.

Let $P = \{v_i^1 w_i^1 : 1 \leq i \leq n\}$ be the set of forbidden edges in \mathcal{I}' . The edges connecting V' and E' capture the incidence relations of the original graph G , while vertices in W and F can be seen as garbage collectors.

$$\begin{array}{llll} v_i^1 : & w_i^1 & e(v_i^1) & w_i^2 & (1 \leq i \leq n) & w_i^1 : & v_i^3 & v_i^1 & (1 \leq i \leq n) \\ v_i^2 : & w_i^2 & e(v_i^2) & w_i^3 & (1 \leq i \leq n) & w_i^2 : & v_i^1 & v_i^2 & (1 \leq i \leq n) \\ v_i^3 : & w_i^3 & e(v_i^3) & w_i^1 & (1 \leq i \leq n) & w_i^3 : & v_i^2 & v_i^3 & (1 \leq i \leq n) \\ e_j^1 : & e_j^2 & v(e_j^1) & f_j^2 & (1 \leq j \leq m) & f_j^1 : & f_j^2 & e_j^2 & (1 \leq j \leq m) \\ e_j^2 : & f_j^1 & v(e_j^2) & e_j^1 & (1 \leq j \leq m) & f_j^2 : & e_j^1 & f_j^1 & (1 \leq j \leq m) \end{array}$$

Figure 8: Preference lists in the constructed instance of SR MIN FORBIDDEN.

Finally we define some further notation in \mathcal{I} . For each i , $1 \leq i \leq n$,

let $V_i^c = \{v_i^r w_i^r : r \in \{1, 2, 3\}\}$ and let $V_i^u = \{v_i^r w_i^{r+1} : r \in \{1, 2, 3\}\}$, where addition is taken modulo 3. Note that each V_i^c contains exactly one forbidden edge, while V_i^u has no forbidden edge. Similarly for each j , $1 \leq j \leq m$, let $E_j^1 = \{e_j^1 e_j^2, f_j^1 f_j^2\}$ and let $E_j^2 = \{e_j^1 f_j^2, e_j^2 f_j^1\}$.

Claim 4.9. \mathcal{I}' admits a stable matching in which every vertex is matched.

Proof. Let $M = \bigcup_{i=1}^n V_i^c \cup \bigcup_{j=1}^m E_j^1$. Starting with the argument that each $v_i^r \in V'$, each $e_j^1 \in E'$ and each $f_j^1 \in F$ receive its first-choice partner in M , it is straightforward to verify that M is stable: the remaining vertices prefer, to their partners, only vertices that already have their first-choice partners. Theorem 1.1 implies then that every stable matching in \mathcal{I}' matches every vertex in \mathcal{I}' . ■

In Claims 4.10 and 4.11 we show that G has a vertex cover C where $|C| \leq K$ if and only if \mathcal{I}' has a stable matching M where $|M \cap P| \leq K$.

Claim 4.10. If G has a vertex cover C such that $|C| \leq K$ in \mathcal{I} , then there is a stable matching M in \mathcal{I}' with $|M \cap P| \leq K$.

Proof. We construct a matching M in \mathcal{I} as follows. For each i ($1 \leq i \leq n$), if $v_i \in C$, add V_i^c to M , otherwise add V_i^u to M . For each j ($1 \leq j \leq m$), if $v_{j,1} \in C$, add E_j^2 to M , otherwise add E_j^1 to M . Then $|M \cap P| = |C| \leq K$.

Now we verify that M is stable in \mathcal{I} . Suppose firstly that some $v_i^r \in V'$ has its third-choice partner in M , and prefers $e(v_i^r)$. Then $v_i \notin C$. Let $e_j^s = e(v_i^r)$. Then by definition, $e_j = e_{i,r}$ and $v_i = v_{j,s}$. Since $v_{j,s} \notin C$, e_j is covered by C at its other endpoint (i.e., $v_{j,3-s}$). By construction of M , $E_j^s \subseteq M$. Hence e_j^s has its first-choice partner in M .

Now suppose that some e_j^s has its third-choice partner in M , and prefers $v(e_j^s)$. Then $E_j^{3-s} \subseteq M$. It follows that $v_{j,s} \in C$, since C is a vertex cover. Let $v_i^r = v(e_j^s)$. Then by definition, $e_j = e_{i,r}$ and $v_i = v_{j,s}$. Thus $v_i \in C$. By construction of M , $V_i^c \subseteq M$. Hence v_i^r has its first-choice partner in M . ■

Claim 4.11. If there is a stable matching M with $|M \cap P| \leq K$ in \mathcal{I}' , then G has a vertex cover C in \mathcal{I} such that $|C| \leq K$.

Proof. We construct a set of vertices C in G as follows. Claim 4.9 states that M matches every vertex in \mathcal{I} . Hence for each i ($1 \leq i \leq n$), either $V_i^c \subseteq M$ or $V_i^u \subseteq M$. In the former case add v_i to C . As $|M \cap P| \leq K$, it follows that $|C| \leq K$. Also, for each j ($1 \leq j \leq m$), as M matches every vertex in \mathcal{I} , either $E_j^1 \subseteq M$ or $E_j^2 \subseteq M$.

Assume that C is not a vertex cover in G , i.e., there is an edge $e_j = v_{j,1}v_{j,2}$ such that $v_{j,1} \notin C$ and $v_{j,2} \notin C$. Suppose that $v_{i_1} = v_{j,1}$ and $v_{i_2} = v_{j,2}$. Then $V_{i_1}^u \subseteq M$ and $V_{i_2}^u \subseteq M$.

Now let $r \in \{1, 2, 3\}$ be such that $e_j = e_{i_1, r}$ and let $r' \in \{1, 2, 3\}$ be such that $e_j = e_{i_2, r'}$. Then $e(v_{i_1}^r) = e_j^1$ and $v_{i_1}^r$ prefers e_j^1 to its partner in M . Similarly $e(v_{i_2}^{r'}) = e_j^2$ and $v_{i_2}^{r'}$ prefers e_j^2 to its partner in M .

If $E_j^1 \subseteq M$ then $v_{i_2}^{r'}e_j^2$ blocks M . Otherwise $E_j^2 \subseteq M$ and $v_{i_1}^re_j^1$ blocks M . This contradiction to the stability of M implies that C is a vertex cover in G . ■

For SR MAX FORCED, an analogous proof can be derived if we define the set of forced edges to be $Q = \{v_i^1w_i^2 : 1 \leq i \leq n\}$. □

Theorem 4.12. SR MIN RESTRICTED VIOLATIONS is solvable in $O(n)$ time if every preference list is of length at most 2.

Proof. Since the the set of matched vertices is the same in all stable matchings by Theorem 1.1, finding a stable matching in $O(n)$ time in these very strongly restricted instances marks all vertices that need to be matched. In each component C , since C is a path or a cycle, there are at most two possible stable matchings satisfying these constraints. We choose the stable matching in C that violates fewer constraints. □

Short preference lists are not the only case when SR MIN RESTRICTED VIOLATIONS becomes tractable, as our last theorem shows.

Theorem 4.13. SR MIN RESTRICTED VIOLATIONS is solvable in polynomial time if the number of restricted edges or the minimal number of violated constraints is constant.

Proof. Suppose firstly that $L = |P| + |Q|$ is a constant. We will show how to solve SR MIN RESTRICTED VIOLATIONS DEC in polynomial time. We assume that, for the purposes of this proof, the problem definition is modified so that, given an instance $\mathcal{I} = (G, O, P, Q, K)$, we are required to find a stable matching M in G such that $|M \cap P| + |Q \setminus M| \leq K$, or report that no such matching exists.

Our first observation is that this problem is trivially solvable if the target value K satisfies $K \geq L$. In this case, any stable matching will suffice.

Now assume that $K < L$. Suppose firstly that there is a stable matching M in G violating $k \leq K$ restrictions. Then $|M \cap P| = k_1$ and $|Q \setminus M| = k_2$,

where $k_1 + k_2 = k$. If we let $P' = M \cap P$ and $Q' = Q \setminus M$ then M is a stable matching in \mathcal{I} containing no edge in $(P \setminus P') \cup Q'$ and containing all edges in $P' \cup (Q \setminus Q')$.

Hence to solve SR MIN RESTRICTED VIOLATIONS DEC we generate all subsets S of $P \cup Q$ of size k , for each $k \leq K$. Then we run the algorithm of [10] to determine in $O(m)$ time whether there is a stable matching containing no edge in $(P \setminus S) \cup (Q \cap S)$ and containing all edges in $(P \cap S) \cup (Q \setminus S)$.

Thus $\sum_{i=0}^K \binom{m}{i} = \sum_{i=0}^L \binom{m}{i}$ subsets are generated to determine whether the desired matching exists. The number of rounds is thus $O(m^L)$, while each round takes $O(m)$ time to complete. The overall running time is $O(m^{L+1})$.

We now show how to use the above approach in order to solve SR MIN RESTRICTED VIOLATIONS. If we find a solution during course of this process then G admits a stable matching M such that $|M \cap P| + |Q \setminus M| \leq K$. In order to minimise $|M \cap P| + |Q \setminus M|$ it suffices to use the above technique in combination with a binary search procedure on values of $K \leq L$. This requires $O(\log L)$ invocations of the algorithm for the decision problem, which is a constant, and hence the overall time complexity remains $O(m^{L+1})$.

Now suppose that L is the minimal number of violated constraints, and that L is constant. For each value of K , where K starts from 0 and increases by 1 after each iteration, we execute the algorithm described above to solve SR MIN RESTRICTED VIOLATIONS DEC, noting that it is sufficient to generate all subsets of size exactly K at each iteration. We terminate as soon as we find a stable matching M such that $|M \cap P| + |Q \setminus M| \leq K$. This process is bound to halt, since by definition, \mathcal{I} admits a stable matching M such that $|M \cap P| + |Q \setminus M| = L$, so $K \leq L$. Thus the overall time complexity of this approach is $O((L+1)m^{L+1}) = O(m^{L+1})$, which is polynomial if L is a constant. \square

5. Conclusion and open questions

In this paper, we investigated the stable marriage and the stable roommates problems on graphs with forced and forbidden edges. Since a solution satisfying all constraints need not exist, two relaxed problems were defined. In MIN BP SM RESTRICTED, constraints on restricted edges are strict, while a matching with the minimum number of blocking edges is searched for. On the other hand, in SR MIN RESTRICTED VIOLATIONS, we seek stable solutions that violate as few constraints on restricted edges as possible. For both problems, we determined the complexity and studied several special cases.

One of the most striking open questions is the approximability of SR MIN RESTRICTED VIOLATIONS. Even though the problem can easily be formulated as a weighted SR problem, the 2-approximation for this latter problem [33, 34] only holds for instances with specific P and Q sets. This is due to the non-negativity and monotonicity constraints on the 2-approximation result. Another open question is formulated as Conjecture 3.18: the complexity of MIN BP SM RESTRICTED is not known if each woman’s preference list consists of at most 2 elements.

A more general direction of further research involves the SM MIN RESTRICTED VIOLATIONS problem. We have shown that it can be solved in polynomial time, using algorithms for minimum weight stable marriage. The following question arises naturally: is there a faster method for SM MIN RESTRICTED VIOLATIONS that avoids reliance on Feder’s algorithm or linear programming methods?

Another natural generalisation is to consider preference lists involving ties. Our hardness results carry over to this case, but the positive results need to be revisited.

Besides the two main problems discussed in this paper, other approximation concepts can also be investigated in the framework of restricted edges. For example one alternative would be to combine the two objectives that we considered. Can we efficiently find matchings that minimise the total number of violated constraints, that is, $|bp(M)| + |M \setminus Q| + |M \cap P|$?

Counting the number of blocking pairs is the most prevalent, but not the only relaxation of stability that has been studied in the literature. Other relaxations can also be combined with the presence of restricted edges in the instance, such as the following concepts.

- *Maximum internally stable matchings* [31], where the goal is to maximise the set of pairs that are stable within themselves.
- *Maximum irreversible stable matchings* [4], where the goal is to maximise the number of irreversible pairs. A pair is *irreversible* if, once it is contained in a matching, no agent from the pair will ever be in a blocking edge, irrespective of how the outside agents are matched.
- *Socially stable matchings* [2], where only a fixed subset of pairs (described by a given social network graph) can be blocking.

Acknowledgements

We would like to thank the anonymous reviewers of this paper and an earlier version of it for their valuable comments, which helped to improve the presentation, and for suggesting several of the open problems that feature in Section 5.

References

- [1] Abraham, D. J., Biró, P., Manlove, D. F., 2006. “Almost stable” matchings in the roommates problem. In: Erlebach, T., Persiano, G. (Eds.), *Proceedings of WAOA '05: the 3rd Workshop on Approximation and Online Algorithms*. Vol. 3879 of *Lecture Notes in Computer Science*. Springer, pp. 1–14.
- [2] Askalidis, G., Immorlica, N., Kwanashie, A., Manlove, D., Pountourakis, E., 2013. Socially stable matchings in the Hospitals / Residents problem. In: *Proceedings of WADS 2013: the 13th Algorithms and Data Structures Symposium*. Vol. *Lecture Notes in Computer Science* of 8037. Springer, pp. 85–96.
- [3] Berman, P., Karpinski, M., Scott, A. D., 2003. Approximation hardness of short symmetric instances of MAX-3SAT, *electronic Colloquium on Computational Complexity Report*, number 49.
- [4] Biró, P., Iñarra, E., Molis, E., 2016. A new solution concept for the roommate problem: Q-stable matchings. *Mathematical Social Sciences* 79, 74–82.
- [5] Biró, P., Manlove, D. F., McDermid, E. J., 2012. “Almost stable” matchings in the roommates problem with bounded preference lists. *Theoretical Computer Science* 432, 10–20.
- [6] Biró, P., Manlove, D. F., Mittal, S., 2010. Size versus stability in the marriage problem. *Theoretical Computer Science* 411, 1828–1841.
- [7] Dias, V. M. F., da Fonseca, G. D., de Figueiredo, C. M. H., Szwarcfiter, J. L., 2003. The stable marriage problem with restricted pairs. *Theoretical Computer Science* 306, 391–405.

- [8] Feder, T., 1992. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences* 45, 233–284.
- [9] Feder, T., 1994. Network flow and 2-satisfiability. *Algorithmica* 11, 291–319.
- [10] Fleiner, T., Irving, R. W., Manlove, D. F., 2007. Efficient algorithms for generalised stable marriage and roommates problems. *Theoretical Computer Science* 381, 162–176.
- [11] Gai, A.-T., Lebedev, D., Mathieu, F., de Montgolfier, F., Reynier, J., Viennot, L., 2007. Acyclic preference systems in P2P networks. In: Kermarrec, A., Bougé, L., Priol, T. (Eds.), *Proceedings of Euro-Par '07 (European Conference on Parallel and Distributed Computing): the 13th International Euro-Par Conference*. Vol. 4641 of *Lecture Notes in Computer Science*. Springer, pp. 825–834.
- [12] Gale, D., Shapley, L. S., 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–15.
- [13] Gale, D., Sotomayor, M., 1985. Some remarks on the stable matching problem. *Discrete Applied Mathematics* 11, 223–232.
- [14] Garey, M. R., Johnson, D. S., 1979. *Computers and Intractability*. Freeman, San Francisco, CA.
- [15] Garey, M. R., Johnson, D. S., Stockmeyer, L., 1976. Some simplified NP-complete graph problems. *Theoretical Computer Science* 1, 237–267.
- [16] Gusfield, D., Irving, R. W., 1989. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press.
- [17] Hamada, K., Iwama, K., Miyazaki, S., 2009. An improved approximation lower bound for finding almost stable maximum matchings. *Information Processing Letters* 109, 1036–1040.
- [18] Irving, R. W., 1985. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms* 6, 577–595.
- [19] Irving, R. W., Leather, P., Gusfield, D., 1987. An efficient algorithm for the “optimal” stable marriage. *Journal of the ACM* 34, 532–543.

- [20] Irving, R. W., Manlove, D. F., 2002. The stable roommates problem with ties. *Journal of Algorithms* 43, 85–105.
- [21] Irving, R. W., Manlove, D. F., O’Malley, G., 2009. Stable marriage with ties and bounded length preference lists. *Journal of Discrete Algorithms* 7 (2), 213–219.
- [22] Khot, S., Regev, O., 2008. Vertex cover might be hard to approximate to within $2 - \varepsilon$. *Journal of Computer and System Sciences* 74, 335–349.
- [23] Khuller, S., Mitchell, S., Vazirani, V., 1994. On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science* 127, 255–267.
- [24] Knuth, D., 1976. Mariages Stables et leurs relations avec d’autres problèmes combinatoires. Les Presses de L’Université de Montréal, English translation in *Stable Marriage and its Relation to Other Combinatorial Problems*, volume 10 of CRM Proceedings and Lecture Notes, American Mathematical Society, 1997.
- [25] Maier, D., Storer, J., October 1977. A note on the complexity of the superstring problem. Tech. Rep. 233, Princeton University, Department of Electrical Engineering and Computer Science, Princeton, NJ.
- [26] O’Malley, G., 2007. Algorithmic aspects of stable matching problems. Ph.D. thesis, University of Glasgow, Department of Computing Science.
- [27] Roth, A. E., 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy* 92, 991–1016.
- [28] Roth, A. E., 2008. Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory* 36, 537–569.
- [29] Roth, A. E., Sotomayor, M. A. O., 1990. Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis. Vol. 18 of *Econometric Society Monographs*. Cambridge University Press.
- [30] Rothblum, U. G., 1992. Characterization of stable matchings as extreme points of a polytope. *Mathematical Programming* 54, 57–67.

- [31] Tan, J., 1990. A maximum stable matching for the roommates problem. *BIT Numerical Mathematics* 29, 631–640.
- [32] Tan, J., 1991. A necessary and sufficient condition for the existence of a complete stable matching. *Journal of Algorithms* 12, 154–178.
- [33] Teo, C.-P., Sethuraman, J., 1997. LP based approach to optimal stable matchings. In: Saks, M. E. (Ed.), *Proceedings of SODA '97: the 8th ACM-SIAM Symposium on Discrete Algorithms*. ACM-SIAM, pp. 710–719.
- [34] Teo, C.-P., Sethuraman, J., 1998. The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research* 23, 874–891.
- [35] Zuckerman, D., 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing* 3, 103–128.